

OSG All Hands Meeting 2011


A Glide-in Factory for OSG VOs to Use

By Jeff Dost (UCSD)

Outline

- **GlideinWMS Overview**
 - Summary of what it is
 - How VO users benefit from it
- **UCSD OSG Factory**
 - How a VO Frontend registers with us
 - Details of the service we provide (What we do)
 - Factory Statistics
 - SLA
 - Change Management Plans

Grid Complexities VO Users Face



- Different job managers – Condor, PBS, SGE, etc.
- rsl strings
- Memory constraints
- Runtime constraints

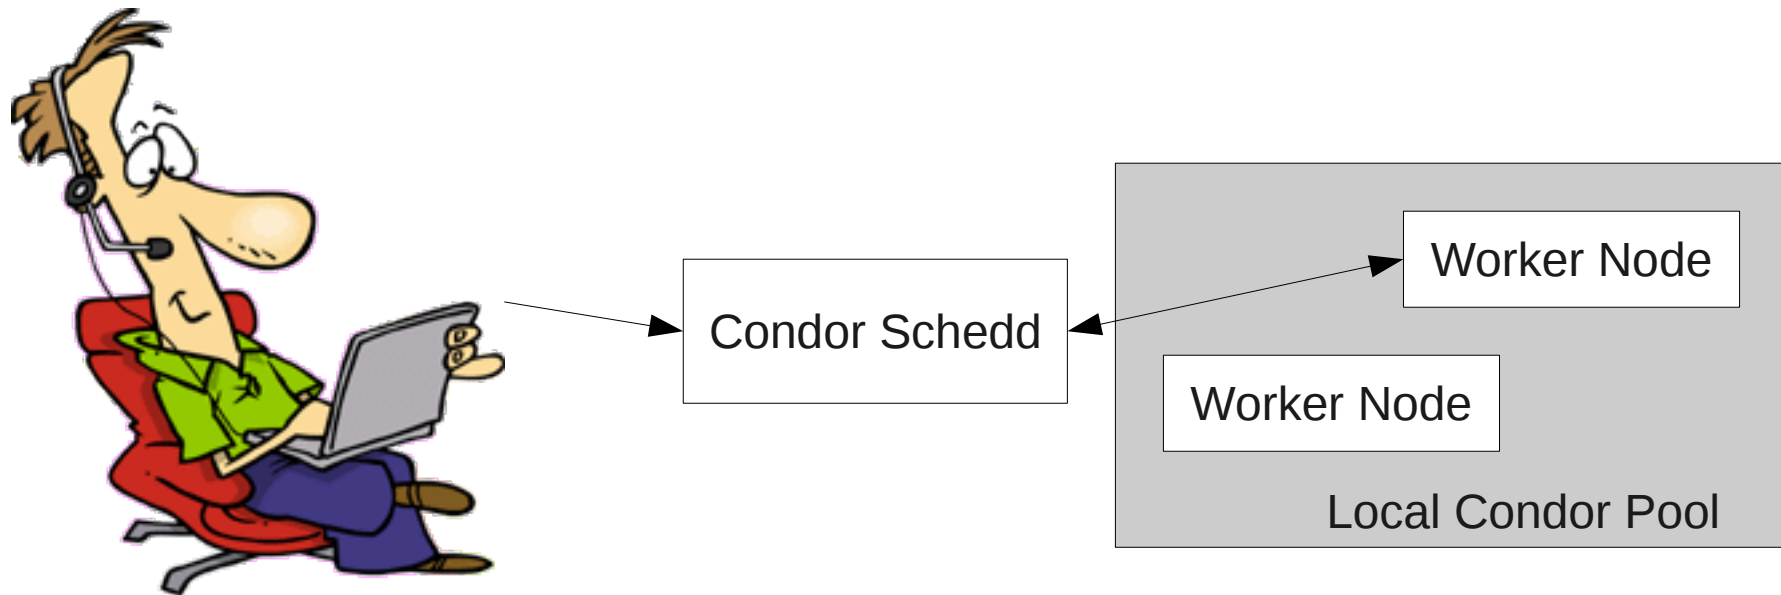
VO users also have to make sure resources are available and must be able to adapt to site changes and failures.

GlideinWMS in a nutshell

GlideinWMS simplifies the life of the VO user by hiding the grid infrastructure and delegating the grid submission duties to the Glidein Factory.

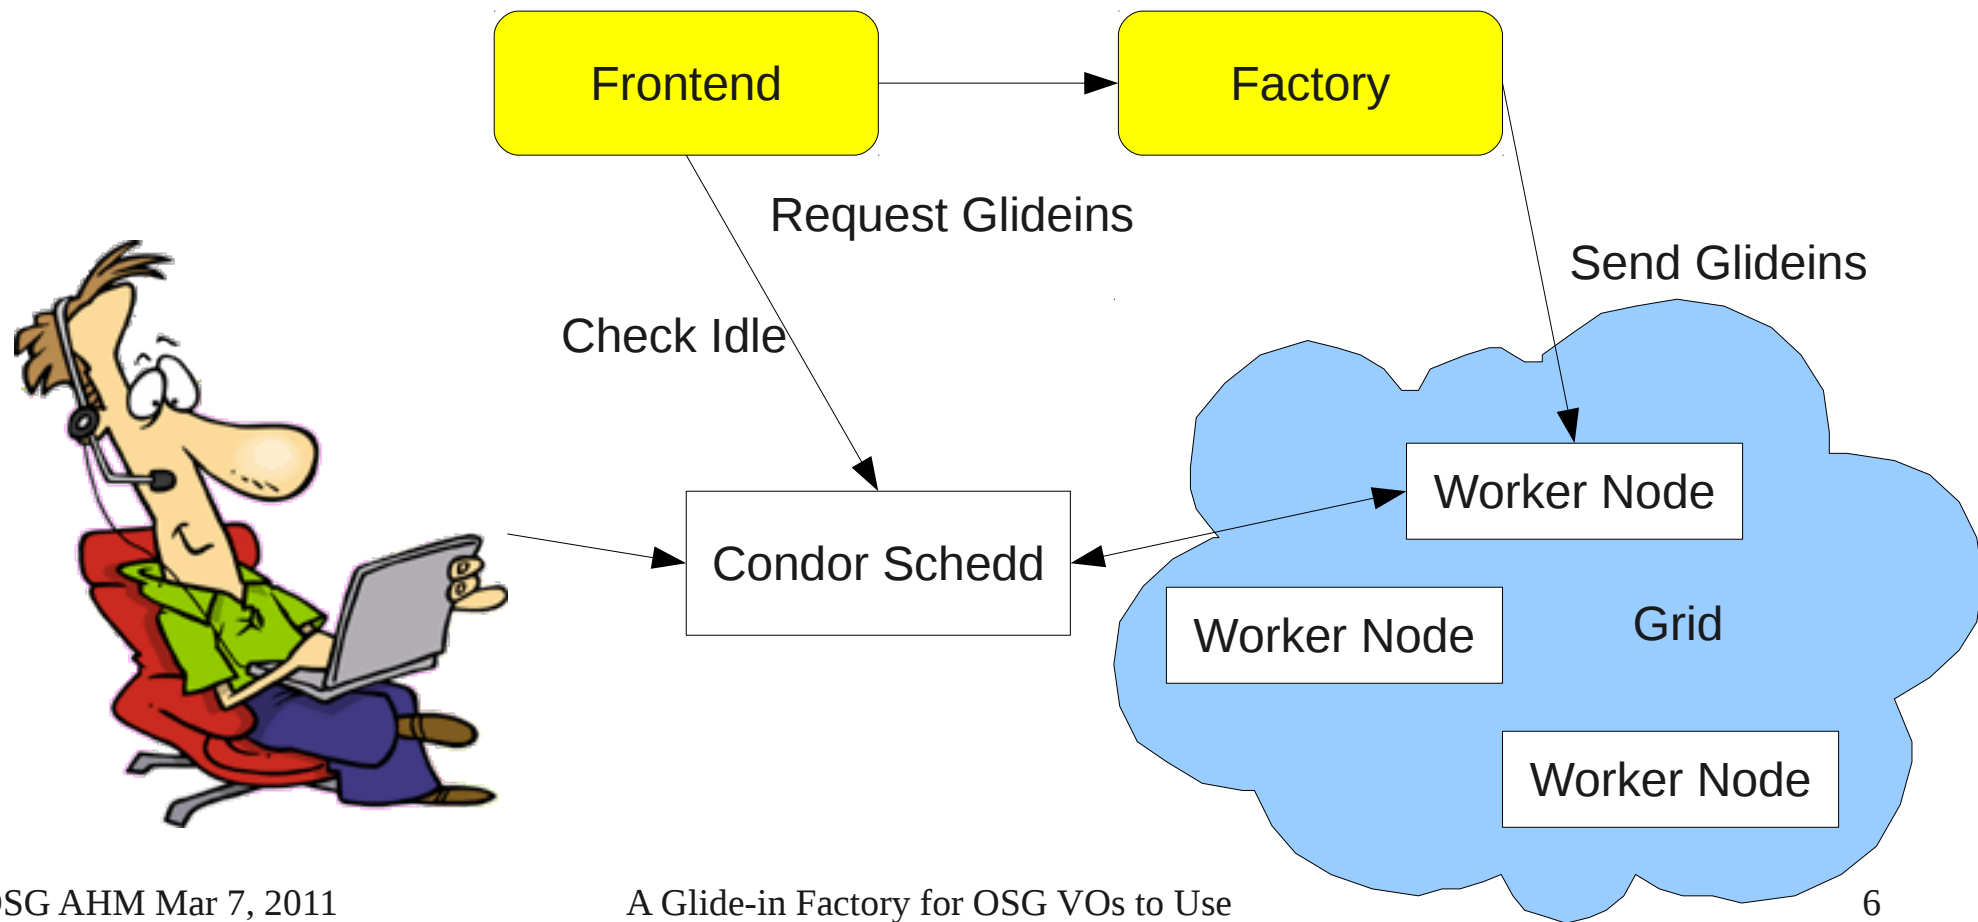
GlideinWMS (a simplified overview)

- Using Condor on local pool



GlideinWMS (a simplified overview)

- Using Condor with GlideinWMS



Benefits of GlideinWMS

- Provides a homogeneous condor layer that runs on top of the grid
- Hides the complexities of grid job submission from the VO user
- Pool size grows and shrinks as needed
- If grid resources fail user jobs will just start at different sites

Glidein Security

- To run securely, glideins must run with gLExec enabled
 - **Condition for using gLExec** - users must* provide proxies in order to use it

* ways to remove this requirement are being worked on

GlideinWMS Frontend

- Any VO wanting to use our Factory must install a GlideinWMS Frontend
- Frontend consists of 2 main pieces:
 - Condor job submission portal for the end user
 - From the user's perspective this is no different than submitting to a normal Condor schedd
 - Glidein Frontend Interface
 - Communicates to factory and requests glideins as needed

Registration with the OSG Factory

- Setup:
 - Supply us a proxy DN for your Frontend
 - Install the Frontend and Condor schedd / collector
 - Give us a list of the grid resources you would like to submit to
 - We do the rest of the heavy lifting

What We Do


- Configure and enable resources for Frontends to use
 - Glidein Factory exists to handle site configuration / management so VO doesn't have to
 - This an area that we can improve on
 - There are still cases where we depend on our Frontend VO expertise for configuration details, specifically with rsl strings.

What We Do

- Site debugging
 - We have a set of monitoring tools to ensure glideins are running as expected
 - If we see something is wrong we open GOC tickets and work closely with the site to debug

RRD file Loaded files 1/1: total/Status_Attributes.rrd

Resolution: 5min (2 days 13h total) ⌵



Select elements to plot:

- Running glidein jobs
- Max requested glideins
- Glideins at Collector
- Glideins claimed by user jobs
- Glideins not matched
- User jobs running
- User jobs idle
- Requested idle glideins
- Idle glidein jobs
- Info age

XML last update: Mon Mar 7 06:19:56 2011

Entry Name		Stat			
		Running	Idle	Waiting	Pending
CMS_T2_US_UCSD_gw2	↑	404	50	0	50
CMS_T2_US_UCSD_gw4	↑	344	66	0	66
CMS_T2_US_Nebraska_Husker	↑	117	93	0	93
HCCHIPC_T2_US_Purdue_Lepton	↑	84	2	0	2
CMS_T2_US_Nebraska_Red	↑	91	73	0	73

What We Do


- Site Validation
 - We do basic validation on sites
 - We need help from VOs
 - We are not perfect - work in progress to make it better
 - Want to better ensure our current configurations are allowing glideins to run optimally on sites
 - We need better ways to test configurations on new sites before adding them as Factory entries
 - Currently developing tools to send test jobs to sites to verify configurations work

UCSD Factory Statistics

- Currently active VO Frontends:
 - CMS Analysis
 - HCC (Nebraska)
 - SBGrid / NeBioGrid
 - CDF
 - GlueX
 - UCSD Campus Grid (under Engage)
 - GLOW (Wisconsin)
 - nanoHUB

UCSD Factory Statistics

- 160 entries total
- 132 CMS sites
 - All OSG CMS T2 sites used
- 45 entries are shared by at least 2 VOs
- 94 European CMS sites




Service License Agreement

- The factory is open for use by every OSG VO
- We guarantee service during normal business hours
- Do our best to provide service outside the regular 8 x 5
- Can be found at:
<https://twiki.grid.iu.edu/bin/view/Operations/GlideInWMSServiceLevelAgreement>

Plans to Improve Change Management

- Currently factory is bleeding edge
 - No obvious problems, but for VOs it is important we minimize the risk
- Coordinating with GOC to improve reliability
- Proposal in the works:
 - **Set up redundant Factory instance** – guarantees high availability
 - **Host second instance at GOC** – provides higher quality hardware to increase dependability

Who runs this thing anyway?



- Run by:
 - Jeff Dost - UCSD Physics Computing Staff (Funded by OSG)
 - Ian MacNeill - UCSD Physics Grad Student
- Under supervision of (and providing expert support):
 - Igor Sfiligoi - UCSD Physics Computing Staff (Funded by OSG)
- And to keep us all in check:
 - Frank Wuerthwein – UCSD Physics Professor (and OSG Executive Team member)